



Learnings from curating a trustworthy, well-annotated, and useful dataset of disordered English speech

Pan-Pan Jiang¹, Jimmy Tobin¹, Katrin Tomanek¹, Robert L. MacDonald¹, Katie Seaver², Richard Cave³, Marilyn Ladewig⁴, Rus Heywood⁴, Jordan R. Green²

¹Google Research ²MGH-IHP ³Language and Cognition UCL ⁴CPUnlimited

pjiang@google.com, jtobin@google.com

Abstract

Project Euphonia, a Google initiative, is dedicated to improving automatic speech recognition (ASR) of disordered speech. A central objective of the project is to create a large, high-quality, and diverse speech corpus. This report describes the project's latest advancements in data collection and annotation methodologies, such as expanding speaker diversity in the database, adding human-reviewed transcript corrections and audio quality tags to 350K (of the 1.2M total) audio recordings, and amassing a comprehensive set of metadata (including more than 40 speech characteristic labels) for over 75% of the speakers in the database. We report on the impact of transcript corrections on our machine-learning (ML) research, inter-rater variability of assessments of disordered speech patterns, and our rationale for gathering speech metadata. We also consider the limitations of using automated off-the-shelf annotation methods for assessing disordered speech.

Index Terms: Disordered speech, automatic speech recognition, dysarthria, speech annotation tools

1. Introduction

The Project Euphonia [1] data corpus is, to our knowledge, the largest dataset of disordered English speech in the world, and currently comprises over 1.2M million utterances recorded from ~2000 speakers with varied etiologies. The long-term goal of the project is to make speech recognition technology accessible to everyone, regardless of their speech abilities, similar to work in [2, 3, 4]. This paper outlines the progress and scale of Project Euphonia's dataset, emphasizing its role in improving speech recognition for people with speech disorders. It details enhancements in data diversity, trustworthiness, and collection efficiency, aiming to support the development of a high-quality speech corpus for accelerating ML-driven research on disordered speech.

2. Expanding Data Diversity

2.1. Increasing the diversity of speakers and etiologies

We increased our effort to enroll speakers from a broad spectrum of backgrounds, including individuals from historically underrepresented racial and socioeconomic groups and covering a wide variety of etiologies of speech disorders. Our approach involved forming strategic alliances with diverse organizations dedicated to serving people with communication impairments such as ADAPT Community Network and LSVT Global. In total, our corpus includes over 2000 speakers with varied etiologies including amyotrophic lateral sclerosis (ALS), Parkinson's disease (PD), Down syndrome (DS), laryngectomy (LAR), vocal cord paralysis (VCP), hearing impairment (HI),

multiple sclerosis (MS), cerebral palsy (CP). In addition, to make our data collection efforts more inclusive, we tailored our strategy to cater to participants with different literacy levels. This included the creation of a new set of prompts, composed of shorter phrases and utilizing vocabulary suitable for a Grade 3 to Grade 5 reading level, as measured by the Flesch-Kincaid scale [5].

2.2. Documenting diversity of atypical speech patterns

To ascertain the diversity of the types of atypical speech patterns in our dataset, we developed a human-centric labeling system that could effectively capture a wide range of speech and voice irregularities. This method involved the use of evaluations by certified speech-language pathologists (SLPs), who received specialized training in assessing and grading speech abnormalities. Given the substantial heterogeneity of speech impairments, we graded each sample based on 40 different disordered speech labels. The grading scheme was consistent across the 40 different labels, with most rated along a five-point equal interval scale: typical (no impairment), mild impairment, moderate impairment, severe impairment, and profound impairment. These labels were generated for each of the speech subsystems: articulatory, phonatory, resonatory, and respiratory, and for global measures of speech such as overall severity, speech intelligibility, speaking rate, speech pattern consistency, and prosody. Additional labels were evaluated for the phonatory system including vocal pitch and loudness, and dysphonia characteristics, which included roughness, breathiness, strain, and tremor.

While tools such as openSMILE [6] provide options for automatically labeling disordered speech and voice, our choice to employ expert human raters was driven by the ongoing debate over the clinical validity of quantitative acoustic metrics [7, 8]. The substantial operational costs of obtaining expert human labels, however, necessitated the development of a brief yet comprehensive labeling system, enabling the utilization of a single rater instead of relying on multiple raters or reaching a consensus. In total, we amassed a comprehensive set of these annotations for over 75% of the speakers in our database.

2.3. Expanding linguistic and speech pattern diversity

To enhance the linguistic diversity of our dataset, we transitioned from having all participants record identical phrases to a system where each participant recorded a few phrases randomly selected from a broader pool. We also extended our data collection efforts to include spontaneous speech samples. Unlike the more controlled, prompted speech, spontaneous speech is produced with less precision [9] and may be less intelligible [10, 11]. To collect spontaneous speech samples, we initiated a Trusted Tester Program, inviting a select group of participants

who were interested in contributing to our research and willing to provide additional consent. These datasets consisted of both speech elicited using assessor-prompted conversation cues and fully self-generated speech. Before integrating these samples into our research corpus, we conducted a thorough review to identify and remove any Personally Identifiable Information (PII), ensuring the privacy and security of the data. The samples were added to our research corpus only after this careful screening process.

3. Improving trustworthiness of corpus

3.1. Manual data validation and cleaning

As our corpus of disordered speech grew in size, human intervention was required to facilitate the detection and removal of low-quality data, such as recordings with poor audio quality, empty recordings, and recordings made by a bot or spammer. We also standardized transcriptions to normalize all prompts to follow the same transcription guidelines (e.g., “6 am” vs “6:00 a.m.”). Lastly, we created human-reviewed and corrected transcripts for 352,130 utterances, which accounts for 29% of our overall corpus, with a focus on reviewing test utterances (> 40% reviewed at the time of writing).

We ran experiments to understand the benefit of transcript correction on model accuracy. We chose to focus our assessments on speakers with Down Syndrome since these speakers tended to require more transcript corrections (usually due to reading errors), thus affording us an opportunity to assess any effects those corrections may have. In the end, from a pool of 1000 speakers, we were able to find nine speakers with Down Syndrome (all rated as having an overall mild severity of speech impairment) for whom we had an adequate number of recordings (at least 350), and who had at least 15% of their transcripts reviewed, which in this case resulted in between 700 and 3800 manually corrected transcripts.

On average we standardized 72% of each speaker’s transcripts, with the vast majority of corrections (95% on average) due to punctuation, capitalization, and other normalization aspects of the transcript. Actual word changes (i.e., misreadings) occurred on < 5% of the utterances on average. While many of these transcription errors may not have been meaning-changing, they did impact reported WERs.

We analyzed the impact of corrected transcripts on the training process. We trained models in a similar fashion to [2], with and without corrected transcripts. After evaluating on corrected test transcripts, results from training with corrected transcripts were inconclusive. For some speakers, we saw improvements; for others, we saw small deteriorations (overall, there was no statistically significant improvement across the 9 speakers).

We then analyzed the impact of corrected transcripts on model quality assessment (evaluation). WERs decreased when using corrected transcripts (compared to using uncorrected transcripts) in 7 of the 9 speakers (with deltas ranging from -0.2 WER to -4.5 WER). And increased in 2 of the 9 speakers (deltas of +0.3 WER and +0.9 WER). This meant that most of the time, assessing WERs on uncorrected transcripts led to an overestimation of WERs.

Based on these results, we concluded that transcript normalizations should be consistently applied across all transcripts. Consequently, for Project Euphonia, we automated the application of transcript normalizations for our own data collection operations.

3.2. Testing automated approaches for identifying low-quality data

Manual quality control and labeling are expensive and subjective. Therefore, our group has been exploring a variety of methods to automate quality control. To date, our efforts have primarily focused on developing tools that can automatically detect poor-quality audio recordings. Recording quality can vary from participant to participant and even from recording session to recording session. In some cases, recordings do not contain speech because of user error. Below, we outline our methodology for evaluating a Voice Activity Detection (VAD) strategy aimed at identifying recordings within our disordered speech dataset that, in error, did not contain speech.

VAD models classify a sequence of acoustic frames as either speech or silence [12]. A VAD is part of the process of end-pointing, which can filter and segment speech recordings before the recordings are run through an ASR model. VAD precision is known to be affected by the level of noise in a speech recording, but [13] shows that precision for classifying typical speech reduces from 90% to 80% when noise is introduced into a signal. We analyzed the performance of two off-the-shelf VAD models (VAD1, a cloud-based speech VAD similar to [14], and VAD2, which was designed for on-device ASR similar to [15]) for detecting voice in our recordings of disordered speech.

VAD1 was run on 825k utterances. ~8k utterances were found to be marked as not containing speech. SLPs listened to 7.5% of the ~8k utterances (596 utterances) to determine the false omission rate (Figure 1). The false omission rate is defined as the number of utterances marked as not containing speech which actually contained speech. Of the 378 total speakers who had utterances without detected speech, 117 speakers (31%) had utterances assessed (Figure 1). We found that VAD1 had a false omission rate of 47.7%. This false omission rate is evidence that the atypical speech patterns have an effect distinct from signal noise alone where we would expect 80% precision. Further, certain etiologies were associated with the highest false omission rates including LAR, VCP, MS, and DS.

Speech systems that employ VADs may unintentionally worsen the user experience for individuals with speech disabilities. Specifically, within numerous speech applications, the VAD might exclude certain audio samples from transcription.

To simulate the effect of using a VAD to automatically reject speech recordings from a disordered speech corpus, we selected a subset of 20 speakers with etiologies mentioned above as being most affected by VAD bias and compared the performance of VAD1 to VAD2 on their 32,000 speech recordings. Based on these simulations, we found that a pipeline with VAD1 would reject ~4,000 recordings, while one with VAD2 would reject ~500 recordings. Therefore, the use of VAD1 instead of VAD2 would have resulted in the rejection of on average 175 utterances for each speaker (which is equivalent to an hour or more of user effort). This difference is evidence that not all VAD models are created equal and care should be taken when they are integrated into data curation pipelines. A more permissive VAD is one option, but a VAD trained with impaired speech is very likely a better solution. As a result, we do not use VAD models in any part of our data curation pipeline.

3.3. Establishing the replicability of human expert labeling

Reliable and accurate labeling is essential in ML research, providing the foundation for developing models that are effective and trustworthy. We tested the inter-rater reliability of our disordered speech labels by having two SLPs independently rate




Figure 1: Analyzed VAD negative decisions grouped by etiology. False omission rate is labeled for each etiology. (Top) utterance counts, (bottom) speaker counts.

50 recordings (Figure 2). Krippendorff’s alpha [16] showed that reliability can vary significantly across labels. Although most of the labels achieved substantial or moderate agreement, the consistency and speaking rate labels were unreliable. In contrast, the severity label was highly reliable, a finding that is consistent with prior reports [17]. Taken together, these findings support the need to establish the reliability of disordered speech labels prior to deploying them for the purpose of model building.

4. Improving Data Collection Efficiency

Prior research has shown that even small amounts of speech recordings (20 minutes) can be used to personalize on-device automatic speech recognition (ASR) models for use in home automation and caregiver situations [2]. Based on this finding, we reduced the number of phrases assigned per user from 1500 to 300. To further improve data collection efficiency, we curated our prompt lists to eliminate prompts that were inappropriate for our users (e.g., too much emphasis on physical mobility) or too difficult to pronounce. We also better-aligned prompts with use cases of interest for our target populations (e.g., home automation, voice search queries).

5. Data Access

The absence of large, well-curated speech-disordered databases has posed a substantial hurdle to advancements in our field. While the primary aim of this paper is to detail the processes, procedures, and underlying rationale for curating this novel dataset, we understand researchers will be most interested in the dataset itself, the value of which tran-




Figure 2: Estimates of inter-rater reliability for nine of the disordered speech labels.

sends academic and industry boundaries. As such, we helped establish the Speech Accessibility Project (SAP - <https://speechaccessibilityproject.beckman.illinois.edu>) - a collaboration between researchers at the University of Illinois Urbana-Champaign (UIUC), and five technology companies, including Google. The SAP aims to collect and curate datasets of impaired speech that will be made available to requestors who sign the UIUC’s Data Use Agreement and whose application is deemed aligned with the program’s objectives by UIUC. Project Euphonia regularly provides advisory support to the SAP based in large part on learnings represented in this paper. Most importantly: Project Euphonia’s 2000 participants, whose contributions comprise the 1.2M utterance dataset herein described, can choose to share their recordings with the SAP so that their data can be studied more broadly.

6. Summary and Future Directions

In this report, we described the iterative improvements made to our dataset of disordered English speech, including updates to the speech sets participants are asked to record; manual edits of recorded audio to create groundtruth transcripts; an adoption of annotation systems developed to create useful metadata and labels; and research into the limitations of some off-the-shelf automated speech label technologies for disordered speech. The following improvements to our annotation processes and datasets are also now being considered by the team:

- Exploring the efficacy of potentially more granular and reliable speech labeling approaches, such as visual analog scales (VAS), direct magnitude scaling (DMS), or consensus-based ratings. [18].
- Improving automated annotation technologies, which have to date been trained only on typical speech samples, by incorporating disordered speech datasets.
- Expanding our models to include Indic, Asian, and European languages.

These improvements will enhance ASR by providing more accurate and diverse labeling, training the technology with a broader range of speech patterns, and expanding linguistic coverage, ensuring a more inclusive and precise recognition across different languages and speech disorders.

7. References

- [1] R. L. MacDonald, P.-P. Jiang, J. Cattiau, R. Heywood, R. Cave, K. Seaver, M. Ladewig, J. Tobin, M. Brenner, P. Q. Nelson, J. R. Green, and K. Tomanek, "Disordered speech data collection: Lessons learned at 1 million utterances from project euphonia," in *Interspeech*, 2021.
- [2] J. Tobin and K. Tomanek, "Personalized automatic speech recognition trained on small disordered speech datasets," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2022, pp. 6637–6641.
- [3] C. Lea, Z. Huang, J. Narain, L. Tooley, D. Yee, T. D. Tran, P. Georgiou, J. Bigham, and L. Findlater, "From user perceptions to technical improvement: Enabling people who stutter to better use speech recognition," in *CHI*, 2023. [Online]. Available: <https://arxiv.org/abs/2302.09044>
- [4] J. R. Green, R. L. MacDonald, P.-P. Jiang, J. Cattiau, R. Heywood, R. Cave, K. Seaver, M. A. Ladewig, J. Tobin, M. P. Brenner, P. Nelson, and K. Tomanek, "Automatic speech recognition of disordered speech: Personalized models outperforming human listeners on short phrases," in *Interspeech*, 2021, pp. 4778–4782.
- [5] R. Flesch, "A new readability yardstick." *Journal of applied psychology*, vol. 32, no. 3, p. 221, 1948.
- [6] F. Eyben, M. Wöllmer, and B. Schuller, "Opensmile: the munich versatile and fast open-source audio feature extractor," in *Proceedings of the 18th ACM international conference on Multimedia*, 2010, pp. 1459–1462.
- [7] G. M. Stegmann, S. Hahn, J. Liss, J. Shefner, S. B. Rutkove, K. Kawabata, S. Bhandari, K. Shelton, C. J. Duncan, and V. Berisha, "Repeatability of commonly used speech and language features for clinical applications," *Digital biomarkers*, vol. 4, no. 3, pp. 109–122, 2020.
- [8] A. P. Vogel, J. Fletcher, P. J. Snyder, A. Fredrickson, and P. Maruff, "Reliability, stability, and sensitivity to change and impairment in acoustic measures of timing and frequency," *Journal of Voice*, vol. 25, no. 2, pp. 137–149, 2011.
- [9] C. Kuo and G. Weismer, "Vowel reduction across tasks for male speakers of american english," *The Journal of the Acoustical Society of America*, vol. 140, no. 1, pp. 369–383, 2016.
- [10] D. Kempler and D. Van Lancker, "Effect of speech task on intelligibility in dysarthria: A case study of parkinson's disease," *Brain and language*, vol. 80, no. 3, pp. 449–464, 2002.
- [11] P. Weir-Mayta, K. A. Spencer, T. L. Eadie, K. Yorkston, S. Savaglio, and C. Woollcott, "Internally versus externally cued speech in parkinson's disease and cerebellar disease," *American Journal of Speech-Language Pathology*, vol. 26, no. 2S, pp. 583–595, 2017.
- [12] J. Sohn, N. S. Kim, and W. Sung, "A statistical model-based voice activity detection," *IEEE Signal Processing Letters*, vol. 6, no. 1, pp. 1–3, 1999.
- [13] S. Ding, Q. Wang, S.-Y. Chang, L. Wan, and I. Lopez Moreno, "Personal VAD: Speaker-Conditioned Voice Activity Detection," in *Proc. The Speaker and Language Recognition Workshop (Odyssey 2020)*, 2020, pp. 433–439.
- [14] S.-Y. Chang, B. Li, and G. Simko, "A unified endpointer using multitask and multidomain training," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019, pp. 100–106.
- [15] T. N. Sainath, Y. He, B. Li, A. Narayanan, R. Pang, A. Bruguier, S.-y. Chang, W. Li, R. Alvarez, Z. Chen *et al.*, "A streaming on-device end-to-end model surpassing server-side conventional model quality and latency," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2020, pp. 6059–6063.
- [16] J. Hughes, "krippendorffsalpha: An r package for measuring agreement using krippendorff's alpha coefficient," *arXiv preprint arXiv:2103.12170*, 2021.
- [17] K. L. Stipanovic, K. M. Palmer, H. P. Rowe, Y. Yunusova, J. D. Berry, and J. R. Green, "'you say severe, i say mild': Toward an empirical classification of dysarthria severity," *Journal of Speech, Language, and Hearing Research*, vol. 64, no. 12, pp. 4718–4735, 2021.
- [18] N. Schiavetti *et al.*, "Scaling procedures for the measurement of speech intelligibility," *Intelligibility in speech disorders*, pp. 11–34, 1992.